**WHAT IS CLAIMED IS:**

1. A method for exclusive two-level caching in a chip-multiprocessor; comprising:

relaxing the inclusion requirement in a two-level cache system in order to form an exclusive cache hierarchy;

providing a first tag-state structure in a first level cache of the two-level cache system, the first tag-state structure having state information;

maintaining in a second-level cache of the two-level cache system a duplicate of the first tag-state structure;

extending the state information in the duplicate of the first tag-state structure, but not in the first tag-state structure, to include an owner indication;

providing in the second-level cache a second tag-state structure so that a simultaneous lookup at the duplicate of the first tag-state structure and the second tag-state structure is possible; and

at any given time of a cache line lifetime in the chip-multiprocessor, associating a single owner with the cache line.

2. A method as in claim 1, wherein the tag-state structure in the second-level cache includes valid and owner indications.

3. A method as in claim 1, wherein each first-level cache is associated with and is private to a particular processor in the chip-multiprocessor, and wherein all processors in the chip multiprocessor share the second-level cache.

4. A method as in claim 1, wherein the exclusive cache hierarchy is formed to minimize data replication and on-chip data traffic without incurring increased second-level hit latency or occupancy.

5. A method as in claim 1, wherein the duplicate of the first tag-state structure is maintained in the second-level cache in order to allow the simultaneous lookup without having to replicate cache lines.

6. A method as in claim 1, further comprising:

determining which instance of one or more copies of a particular cache is the owner copy of that cache line by using information obtained from the simultaneous lookup respecting that cache line.

7. A method as in claim 1, wherein the simultaneous lookup respecting a particular cache line yields the state information of one or more instances of the particular cache line that are present in the first-level cache and/or the second-level cache, the state information for each instance of that cache line including its owner state, and wherein the owner state of only one instance is owner.

8. A method as in claim 1, wherein if the substantially simultaneous lookup yield no instance of the particular cache line in the first-level cache which is owner, the instance in second-level cache is owner by default.

9. A method as in claim 1, wherein associating a single owner with each cache line eliminates unnecessary second-level fills, thereby maximizing the effective use of the two-level cache in the chip-multiprocessor.

10. A method as in claim 1, wherein the state information in the first tag-state and duplicate of the first tag-state structures includes valid/invalid indication and shared/exclusive indication, wherein a cache line instance found to be invalid cannot be owner, and wherein a cache line instance that is found to be exclusive cannot be involved in a write-back to the second level cache.

11. A method for maximizing the use of on-chip cache memory capacity in a chip-multiprocessor, comprising:

forming a two-level cache system with an exclusive cache hierarchy in order to minimize cache line replication and on-chip traffic, the two-level cache system including a first-level cache dedicated to each processor in the chip-multiprocessor and a second-level cache shared by all the processors;

associating cache lines with an indication of ownership so that among one or more than one instances of each cache line present in the two-level cache system, there is only one instance that is an owner instance, the indication of ownership being provided only in a second-level cache;

associating the cache lines with state information that includes the indication of ownership, the state information for cache lines present in the first-level cache being maintained in the second-level cache; and

administering cache line ownership and write-backs based on a predetermined guideline.

12. A method as in claim 11, wherein based the predetermined guideline a first-level cache miss that finds no other copy of a requested cache line becomes an owner of the cache line.

13. A method as in claim 11, wherein based the predetermined guideline a first-level cache miss that does not find a copy of a cache line in the second-level cache but finds it in one or more than one of the first-level caches receives that cache line from a previous owner and becomes a new owner.

14. A method as in claim 11, wherein based the predetermined guideline a first-level cache that replaces a cache line, is informed by the second-level cache whether it is owner, in which case it issues a second level cache fill.

15. A method as in claim 11, wherein based the predetermined guideline whenever the second-level cache has a copy of a cache line, the second-level cache is the owner, and

wherein a first-level cache miss that hits in the second-level cache without invalidating it by a write miss does not steal ownership from the second-level cache.

16. A method as in claim 11, wherein based the predetermined guideline whenever the second-level cache needs to evict a cache line that is additionally present in one or more first-level caches the second-level cache arbitrarily selects one of these first-level caches as the new owner.

17. A two-level cache system in a chip-multiprocessor; comprising:

means for relaxing the inclusion requirement in the two-level cache system in order to form exclusive two-level caching;

means for providing a first tag-state structure in a first level cache of the two-level cache system, the first tag-state structure having state information;

means for maintaining in a second-level cache of the two-level cache system a duplicate of the first tag-state structure;

means for extending the state information in the duplicate of the first tag-state structure, but not in the first tag-state structure, to include an owner indication;

means for providing in the second-level cache a second tag-state structure so that a simultaneous lookup at the duplicate of the first tag-state structure and the second tag-state structure is possible; and

means for associating a single owner with a cache line at any given time of its lifetime in the chip-multiprocessor.

18. A two-level cache system as in claim 17, wherein the tag-state structure in the second-level cache includes valid and owner indications.

19. A two-level cache system as in claim 17, wherein each first-level cache is associated with and is private to a particular processor in the chip multiprocessor, and wherein all processors in the chip multiprocessor share the second-level cache.

20. A two-level cache as in claim 17, wherein the system can be implemented in hardware or software or a combination of both.

21. A two-level cache system in a chip-multiprocessor for maximizing the use of on-chip cache memory capacity in the chip multiprocessor, comprising:

means for forming a two-level cache system with an exclusive cache hierarchy in order to minimize cache line replication and on-chip traffic, the two-level cache system including a first-level cache dedicated to each processor in the chip-multiprocessor and a second-level cache shared by all the processors;

means for associating cache lines with an indication of ownership so that among one or more than one instances of each cache line present in the two-level cache system, there is only one instance that is an owner instance, the indication of ownership being provided only in a second-level cache;

means for associating the cache lines with state information that includes the indication of ownership, the state information for cache lines present in the first-level cache being maintained in the second-level cache; and

means for administering cache line ownership and write-backs based on a predetermined guideline.

22. A two-level cache as in claim 21, wherein the system can be implemented in hardware or software or a combination of both.

23. A two-level cache system in a chip-multiprocessor, comprising:

a plurality of per-processor first level caches, each including an instruction cache and a data cache and each maintaining a first tag-state structure with state information;

an interconnect device;

a second-level cache shared by all the processors, the per-processor first-level caches interfacing with each other and the second cache via the interconnect device, the second-level cache including one or more modules each of which including,

storage for a second tag-state structure with its associated state information,

a memory controller configured to operatively interface with the processors,

a memory, and

storage for a duplicate of the first tag-state structures, the state information in the duplicate being extended to include an indication of owner,

wherein the two-level cache system is configured so that during a lifetime of a cache line in the chip-multiprocessor only one instance of the cache line, either in the first-level cache or the second-level cache, can be the owner.

24. A two-level cache system as in claim 23, wherein each of the plurality of per-processor first-level cache is configured as an N-way set associative cache, where N is equal to or greater than 1.

25. A two-level cache system as in claim 23, wherein each of the plurality of per-processor first-level cache is configured as an 2-way set associative cache.

26. A two-level cache system as in claim 23, wherein the memory in the second-level cache includes a plurality of DRAM (dynamic read only memory) components.